Supporting Information

Paul J. Ferraro [*], Merlin M. Hanauer [*], and Katharine R.E. Sims [†]

[*]Department of Economics, Andrew Young School of Policy Studies, Georgia State University, PO Box 3992, Atlanta, GA 30302-3992, USA, and [†]Department of Economics and Environmental Studies Program, Amherst College, PO Box 5000, Amherst, MA 01002

Conditions associated with protected area success in conservation and poverty reduction

Data

Costa Rica. For full details on data, see Andam et al. (1,2). Digital layers of protected areas (source: National System of Conservation Area Office, Ministry of Environment and Energy, 2006) were provided by the Earth Observation Systems Laboratory, University of Alberta. Other GIS layers are land use capacity (source: Ministry of Agriculture) and roads digitized from hard copy maps for 1969 (source: Instituto Geográfico Nacional, Ministerio Obras Publicas y Transporte). Summary statistics of the data are presented in Table S1.

For the deforestation analyses, digital forest cover layers are created from either aerial photographs (baseline, 1960) or Landsat Thematic Mapper satellite images.[1] The units of analysis are land parcel pixels that measure three hectares, the minimum mappable area. We use Andam et al.'s (1) data set, which comprises 20,000 randomly selected land parcels that were forested (80% or more canopy cover) in 1960. Forested parcels in a given year receive a value of 1; deforested parcels receive a value of 0. The outcome of interest is change in forest cover between 1960 and 1997. Given all sample parcels were forested in 1960, the outcome measure equals 0 if the parcel is still forested in 1997 and 1 if it is deforested.

Spatial layers of protected status (IUCN categories Ia, I, II, IV and VI are used in the analyses) and other geographic characteristics are used to create a set of covariates for each land parcel (Table S2). For various reasons (e.g., cloud cover), 4,737 land parcels are dropped prior to analysis, leaving 15,283 land parcels, of which, 2,809 were protected prior to 1980. We remove parcels that were protected after 1980 (2,183), leaving 10,291 unprotected land parcels from which matches can be drawn.

For the poverty analyses, data come from the population and housing censuses conducted by the Instituto Nacional de Estadistica y Censos (INEC) in 1973 and 2000. Digitized GIS census segment boundaries for 1973 and 2000 were provided by the Cartography Department at INEC. The unit of analysis is the census tract (segmento). In 1973 Costa Rica contained 4,694 census tracts with an average size of 8.82km (range: 0.00466-836 km). The 1973 census is used as the baseline year and all census data are geocoded to their respective census tracts. Between 1973 and 2000 there was a great deal of segmentation of census tracts, with few of the segmented tracts being proper subsets (or sharing major borders) with the original 1973 census tracts. Through the method of areal interpolation ((2, 3); see below), the 2000 census data are aggregated to fit to the 1973 census tract boundaries so that the data are spatially and temporally comparable.

The poverty measure (poverty index) builds on recent efforts to develop a census-based poverty index for Costa Rica (4), which uses principal components analysis to formulate a temporally comparable index based on variables believed to influence poverty.

Thailand. For full details on data see Andam et al. (2) and Sims (5). Digital layers of protected area boundaries are from the IUCN World Database on Protected Areas (accessed 3/2007; IUCN categories I and II were used in the analyses). Other GIS data and the source layers from which they are derived are slope and elevation (NIMA's Digital Terrain Elevation Data- USGS Global GIS database, 1999); distance to major cities (ESRI World Cities, 2000); distance to roads in 1962 (digitized East Asia Road Map, U.S. Map Service 1964, data from 1962); distance to rail lines, distance to major rivers, proximity to watershed boundaries, distance to mineral deposits, distance to Thai border, and ecoregions (USGS Global GIS database, 1999), average monthly temperature and rainfall (Marc Souris, IRD).

The deforestation analysis is based on two classified layers from 1973 and 2000. The 1973 data are based on Landsat MSS images interpreted by the Tropical Rain Forest Information Center (Michigan State University) and the 2000 data on Landsat TM images interpreted by the Thai Royal Forestry Department (courtesy of Marc Souris). The units of analysis are points which are spaced so as to represent the centroid of a three hectare parcel. The data set is created in a similar manner to the Costa Rica deforestation data set and comprises 20,000 randomly selected points which were forested in 1973. Forested points in a given year receive a value of 1; deforested points receive a value of 0. The outcome of interest is change in forest cover between 1973 and 2000. Given all sample points were forested in 1973, the outcome measure equals 0 if the point is still forested in 2000 and 1 if it is deforested. Spatial layers of protected status and other geographic characteristics are used to create a set of covariates corresponding to each sample point (Table S2).

For the poverty analysis the unit of analysis is a subdistrict (tambon). In descending order of size, Thailand has administrative units of "province," "district," "subdistrict," and "village." The sample consists of subdistricts in the North and Northeast regions, where the majority of protected forest areas are located. We exclude subdistricts that are less than 10 km away from a major city (population > 100,000; all of these cities had been established by the 1960's). The average size of a subdistrict in the sample is 74 sq km; the average population is 5043.

The poverty measure for Thailand (poverty headcount ratio) is the share of the population with consumption below the poverty line. This outcome is derived from a poverty mapping analysis by Healy and Jitsuchon (6), applying the poverty mapping methodology developed by Elbers et al. (7).

Preprocessing

We preprocess the data (8, 9) using matching techniques prior to performing any of the LOESS or PLM analyses. Our primary motivation for matching is not the estimation of an overall average treatment effect on the treated. With the exception of an analysis of Thailand deforestation at the scale used in our study, these impacts have already been estimated (1, 2). We use matching to preprocess the data so that we can estimate conditional average treatment effects on the treated. To ensure that our analyses are as comparable as possible to the studies from which we draw (1, 2), we use the same matching

[1] Earth Observation Systems Laboratory, University of Alberta, Edmonton, AB.

methods to create the same matched data sets as those studies. These methods were chosen in these studies because they generated the best covariate balance.

The key to matching as an identification strategy to estimate average treatment effects on the treated is the balancing of key covariate distributions across treatment arms (protected and unprotected). This covariate balance is achieved in expectation through randomization. Covariate balance is implicit under randomization because each unit of the experimental sample has an equal probability (or more generally, a probability that is known to the experimenter) of being assigned to treatment or control. Therefore, treatment is assigned independent of potential outcomes $Y(1)$ and $Y(0)$ under treatment ($T = 1$) and control ($T = 0$), respectively. In the absence of a treatment, one would expect similar average outcomes from both groups. Similarly, if both groups were to receive (the same) treatment, one would expect similar average outcomes from both groups. In the statistics, epidemiology and social science literature this assumption is termed ignorability of treatment, independence of treatment or unconfoundedness. Stated formally,

$$E[Y(1)|T = 1)] = E[Y(1)|T = 0)] = E[Y(1)] \qquad [1]$$
$$E[Y(0)|T = 1)] = E[Y(0)|T = 0)] = E[Y(0)]\,. \qquad [2]$$

In words, **[1]** simply states that average potential outcome for the treatment group under treatment, $E[Y(1)|T = 1)]$, is equal to the average potential outcome of the control group *had they been treated*, $E[Y(1)|T = 0)]$. Similarly, **[2]** states that the average potential outcome for the treated group *had they not been treated*, $E[Y(0)|T = 1)]$, is equal to the average potential outcome of the control group in the absence of treatment, $E[Y(0)|T = 0)]$. In **[1]** and **[2]**, the terms $E[Y(1)|T = 0)]$ and $E[Y(0)|T = 1)]$ are termed counterfactual outcomes. The fundamental problem for causal inference (10) is the fact that counterfactual outcomes are not observed. However, with treatment assigned at random (and thus independent of potential outcomes), the average outcome for control units can act as the counterfactual for treatment units, and *vice versa*.

Protected areas in Costa Rica and Thailand were not established randomly. Matching seeks to mimic the identification of randomization by balancing key covariates that jointly determine selection into treatment and outcomes. Balance, conditional on key covariates, leads to conditional ignorability or conditional independence. These more restrictive assumptions can be stated formally as the analogs to **[1]** and **[2]**,

$$E[Y(1)|T = 1, X] = E[Y(1)|T = 0, X)] = E[Y(1)|X] \qquad [3]$$
$$E[Y(0)|T = 1, X] = E[Y(0)|T = 0, X)] = E[Y(0)|X]. \qquad [4]$$

Equations **[3]** and **[4]** state that, conditional on similar covariate distributions across treatment arms, the average outcomes for the matched control units, $E[Y(0)|X, T = 0)]$, can be used as the counterfactual for treatment units, and *vice versa*. In other words, by ensuring that the distributions of key covariates are balanced across treatment and control groups, similar methods to those used in randomized experiments can be used to estimate average treatment effects on matched datasets. We present **[1]**-**[4]** for completeness; however, we focus on the estimation of conditional average treatment effects on the treated, for which only **[2]** and **[4]** are necessary.

By ensuring that units are comparable across treatment and control groups, we make the conditional independence assumption (CIA), which is necessary for causal inference, more defensible (11). We extend the CIA by assuming that if average treatment effect on the treated estimates are unbiased, conditional on balance across key covariates, comparisons of subgroups within these balanced sets are also unbiased. This allows for causal inference to be drawn from the LOESS and PLM analyses.

As mentioned in the main text, matching can only account for heterogeneity in observable covariates. If the selection process and outcomes are systematically determined only by observable characteristics (for which one controls) then a treatment effect estimate derived from a matching algorithm that provides balance will be unbiased and consistent. However, if there are unobservable characteristics that also contribute to determining selection and outcomes, then treatment effect estimates, even for a well balanced matched sample, may be biased. There is no way to formally test the conditional independence assumption, however Andam et al. (1, 2) test the robustness of their estimates (which are derived from the same matched sets used in our study) to unobserved heterogeneity.

Matched Datasets. For the Costa Rica data, we use nearest neighbor Mahalanobis covariate matching with replacement to preprocess the socioeconomic and deforestation data. We use the same algorithm and covariates (Table S1) as Andam et al. (1, 2), and thus our resulting matched datasets are nearly identical to those used in their analyses.[2] The resulting socioeconomic matched set comprises 249 protected (prior to 1980) and unprotected census tracts. The resulting deforestation matched set comprises 2,809 protected (prior to 1980) and unprotected land parcels. See Table S1 for description and summary statistics of the covariates used in each Costa Rica matching specification.

For the Thailand socioeconomic data we use propensity score matching with exact matching on district in order to control for baseline fixed effects associated with poverty. This is the same specification and matched set used in Andam et al. (2010)(2) which comprises 197 protected (prior to 1985) and unprotected subdistricts. For the Thailand deforestation data we use Mahalanobis covariate matching, with exact matching on district, to create a dataset that is similar to the Costa Rica deforestation analysis (see Tables S4 and S5 for estimates of ATT and balancing results). The resulting matched set comprises 2,808 protected (prior to 1985) and unprotected land parcels. See Table S2 for description and summary statistics of the covariates used in each Thailand matching specification.

Thailand Deforestation Analysis. To ensure methodological comparability across countries, we perform a similar deforestation analysis to that of Andam et al. (1) for Thailand. Our primary interest was to create a dataset, comparable to the Costa Rica deforestation dataset, with which to perform the heterogeneity analyses. As a point of departure, however, we perform sample average treatment effect on the treated calculations similar to those done in Andam et al. (1). There are two benefits to this approach. First it offers a comparison to the original Costa Rica deforestation analysis (1). Second, it provides an average treatment effect on the treated (ATT) estimate to which we can contrast our heterogeneity analyses.

In creating our deforestation dataset for Thailand we follow the methodology of Andam et al. (1); see their SI Text), all geoprocessing is done in ArcGIS 9.x. We begin by selecting 20,000 random points, spaced so as to represent 3 ha land

[2] The socioeconomic matched set is identical to the final data set in Andam et al. (2). The deforestation matched sets would be exact, but we use a slightly updated protected areas database resulting in slightly more protected observations. The average treatment effect on the treated estimates, however, are not different between the two datasets. We present the balancing results in Table S3.

parcels, from the areas of Thailand that were forested in 1973, our baseline year. Using spatial overlays, we create indicators for parcels that were protected by 1985 (2,808) and parcels that were protected after 1985 (3,423). The analysis is designed to estimate the impact of protected areas that were established prior to 1985 on deforestation outcomes between 1973 and 2000. Therefore, we remove from the pool of potential controls, any parcel that was protected *after* 1985. As a result, our potential pool of controls comprises 13,609 parcels that were never protected prior to 2000.[3] We run a series of overlay analyses on the remaining parcels to assign a value for each of the covariates listed in upper panel of Table S2.

Using these data, we implement regression bias-adjusted nearest neighbor Mahalanobis matching with replacement (9, 12) to estimate the ATT. Point estimates and balancing results can be found in Tables S3 and S4, respectively.[4] Similar to Andam et al. (1), we find that the naive difference in means overestimates the amount of avoided deforestation attributable to the establishment of protected areas. As noted in the main text, this is a finding that is consistent with the general observation that protected areas tend to be placed on land that is less desirable for agriculture, and therefore less likely to be deforested in the absence of protection. The resulting matched dataset is used for the Thailand deforestation heterogeneity analyses described in the main text.

Locally Weighted Scatterplot Smoothing (LOESS)

Three LOESS estimators (13, 14) are performed for each of the covariates in the heterogeneous response to protection analyses: (1) on the protected units only; (2) on the imputed counterfactual control units only, and; (3) on the difference between protected and counterfactual unprotected units, the Average Treatment Effect on the Treated (ATT).

In LOESS the data of interest are the doubles (Y_i, X_i) representing the outcome and covariate values for observation $i \in \{1, 2, ..., N\}$, where N is the number of observations in the dataset. The data are first ordered according to X such that $X_1 \leq ... \leq X_N$. Beginning with the first observation $(i^* = 1)$ in this ordered set, fitted values $\left(\widehat{Y}\right)$ are predicted via a local quadratic regression

$$\widehat{Y}_{i\in s} = \widehat{\beta}_0 + \widehat{\beta}_1 X_{i\in s} + \widehat{\beta}_2 X^2_{i\in s}, \qquad [5]$$

where the vector $\widehat{\beta}$ is estimated from

$$Y_{i\in s} = \beta_0 + \beta_1 X_{i\in s} + \beta_2 X^2_{i\in s} + \epsilon_i, \qquad [6]$$

and only observations that lie within span (s) are used. The total number of observations used for each imputation is therefore $j = sN$. Moving stepwise through the ordered data set, N local regressions are estimated.

For each of these local regressions all of the j observations are assigned a weight (w_d) using the tricubic function

$$w_d = \begin{cases} \left(1 - |d_i|^3\right)^3 & \text{for } 0 \leq |d_i| < 1 \\ 0 & \text{otherwise} \end{cases}, \qquad [7]$$

where d_i is a cardinal distance ratio

$$d_i = \frac{|X_{i^*} - X_i|}{\max\left(|X_{i^*} - X_i|\right)}. \qquad [8]$$

Here X_{i^*} represents the covariate value of the observation for which we are imputing $\widehat{Y}$. The weight w_d reduces the influence of observations according to their disparity in covariate value as compared to the observation being evaluated. The LOESS estimation moves stepwise repeating **[5]**-**[8]** for each (ith) observation, "re-centering" the span s to include an equal number j observations about the ith observation. The result of these N local regressions is N local fit values $\left(\widehat{Y}_i\right)$ and their corresponding standard errors of the fit which can be used to form confidence intervals about each fit value. This standard LOESS process is run on the protected units for each analysis (dash-dot line in Figures S3-S8).

We extend the LOESS methodology in order to offer comparability to the studies from which we draw (1, 2, 15) by including local bias-adjusted imputation of counterfactual (unprotected) outcomes. This type of method is used in the matching literature (e.g., (9, 12)) to impute counterfactual values by plugging the values of treated unit covariates into the coefficients estimated from a regression of control unit covariates on control unit outcomes. The purpose of this imputation is to reduce post-match bias, in finite samples, due to remaining covariate imbalance. This process is like asking the question, "what would the outcomes of protected units have been in the absence of protection had their covariates influenced their outcomes in the same manner as the units that were not protected?"

Our methodology requires us to modify the LOESS procedure. In order to impute counterfactual outcome values for each treated unit, both protected and unprotected units must be used as inputs for the LOESS. Prior to the ith local estimation outlined in equations **[5]**-**[8]**, a counterfactual value for each protected unit outcome in the span (s) is imputed according to

$$\widetilde{Y}_{i\in s} = Y_{i\in s:T=0} + \widehat{\mu}_0\left(X_{i\in s:T=1}\right) - \widehat{\mu}_0\left(X_{i\in s:T=0}\right), \qquad [9]$$

where T is an indicator of treatment (0 and 1 indicating the unit is unprotected or protected, respectively) and $\widehat{\mu}_0(\cdot)$ represents the predicted values obtained from combining the coefficients from a control group regression, of outcome on covariates, with the respective treated or control covariates (see Tables S1 and S2 for a list of the covariates).[5] In addition to estimating a LOESS curve based on these counterfactual outcomes (dotted line in Figures S3-S8), the counterfactual value $\widetilde{Y}_{i\in s}$ from **[9]** of the observation being evaluated (i^*) is stored in a vector for use in evaluating a LOESS for ATT.[6]

The LOESS curve for ATT is estimated using the difference between actual protected unit outcomes (Y_i) and their respective counterfactual outcomes $\left(\widetilde{Y}_i\right)$ from **[9]**,

$$\left(Y_{i\in s} - \widetilde{Y}_{i\in s}\right) = \beta_0 + \beta_1 X_{i\in s} + \beta_2 X^2_{i\in s} + \epsilon_i, \qquad [10]$$

where the corresponding fits are estimated in a similar manner to **[3]**. The standard error of the fit is used to form the confidence band (red/green shaded area) about the ATT LOESS curve (solid line in Figures 1 and S3-S8).

The span for any LOESS estimator must be chosen so as to balance the bias/variance tradeoff. A relatively small span includes fewer data points and is considered to be more localized and therefore less biased. However, there will be greater variation, *ceteris paribus*, within a small span. Conversely, a relatively large span uses more data and produces smoother curves (less variation) that are considered to be more biased.

[3] Due to incongruence in spatial layers, 160 parcels are dropped prior to analysis.

[4] In addition to the covariates listed in Tables S2 and S4, matching is required to be performed within districts (i.e., exact matching on district ID) to control for regional heterogeneity.

[5] The imputations are calculated by plugging the covariates $X_{i\in s:T=1}$ and $X_{i\in s:T=0}$ into the vector of coefficients from the regression $Y_{i\in s:T=0} = X_{i\in s:T=0}\beta_0 + \varepsilon$ to obtain $\widehat{\mu}_0\left(X_{i\in s:T=1}\right)$ and $\widehat{\mu}_0\left(X_{i\in s:T=0}\right)$, respectively.

[6] Imputations within the LOESS were programmed in R v2.10.1. Code is available from authors upon request.

For each of the LOESS estimators that we implement, we set the span (s) equal to 0.75. We choose this span for all analyses because: (1) after experimenting with many specifications we felt that it captured the important underlying variability with relatively little noise; and (2) we wanted to remain consistent across analyses.

Partial Linear Model (PLM)

Model. For all (moderating) covariates introduced in the Study Design section we use a two-stage semiparametric partial differencing linear model (16, 17). The PLM is advantageous in that it allows us to control, linearly, for a vector of covariates that influence the outcome of interest and then map the outcome as a nonparametric function of the covariate of interest.

The data used in the PLM are the triples (Y_i, X_i, Z_i) where Y is the scalar outcome of interest, X is the scalar covariate for which the nonparametric function will be estimated and Z is a vector of covariates for which we wish to control in our estimation. Our first-stage equation is thus

$$Y_i = Z_i\beta + f(X_i) + \epsilon_i, \quad [11]$$

where β is a vector of coefficients and $f(\cdot)$ is an unknown real function. Our intention is to estimate $f(\cdot)$ net of the effects of Z. In order to achieve the final goal of removing the influence of Z on Y we must first remove the influence of X on Y. In the first stage we begin by ordering the data according to X such that $X_1 \leq ... \leq X_N$ where $i \in \{1, ..., N\}$. Yatchew (16, 17) shows that the influence of X on Y can be removed be taking the (first) difference (in [11]) according to X

$$Y_i - Y_{i-1} = (Z_i - Z_{i-1})\beta + (f(X_i) - f(X_{i-1})) + +\epsilon_i - \epsilon_{i-1}, \quad i = 2, ..., N. \quad [12]$$

Under the assumption that $\partial Y/\partial X$ is bounded by a constant, $(f(X_i) - f(X_{i-1}))$ goes to zero as N increases. Intuitively this assumption implies that, when the data are ordered according to X, the marginal influence of X on Y is zero, so that term can be dropped from the equation. OLS can then be run on [12] to return an estimate of $\widehat{\beta}_{diff}$. Yatchew (16, 17) shows that because $\widehat{\beta}_{diff}$ converges sufficiently quickly to β, $Z_i\widehat{\beta}_{diff}$ can be subtracted from both sides of [11] to obtain

$$\begin{aligned} Y_i - Z_i\widehat{\beta}_{diff} &= Z_i\left(\beta - \widehat{\beta}_{diff}\right) + f(X_i) + \epsilon_i & [13] \\ &\cong f(X_i) + \epsilon_i. & [14] \end{aligned}$$

Denoting $Y_i - Z_i\widehat{\beta}_{diff} = Y_i - \widehat{Y}_{i,diff} = \widetilde{Y}$, the combination of the LHS of [13] and RHS of [14] is equivalent to

$$\widetilde{Y}_i = f(X_i) + \epsilon_i. \quad [15]$$

We are now able to estimate $f(\cdot)$, which is the nonparametric relationship between X and Y, net of the effects of Z. We do so for treatment, control and ATT estimates using the same LOESS estimator described above. Loshkin (19) suggests using LOESS in the second stage and wrote a Stata ado file which performs the estimate. We wrote a similar function for R. The code is available from the authors upon request.

Yatchew (16, 17) noted that although $\widehat{\beta}_{diff}$ is an unbiased estimate of β, due to the differencing, $\widehat{\beta}_{diff}$ is relatively inefficient. However, he provides analytical higher order differencing weights that can be applied to a high order difference generalization of [12] to greatly improve the efficiency of estimates. We incorporate these weights into our estimation using the 10th order difference (the highest order for which weights are provided).[7] See Yatchew (17) for detailed description of the efficiency issues and a table of the analytical weights.

Empirical Specifications. For each of our PLM analyses we include in Z covariates that we believe affect the outcome of interest. This means that we control for the covariates used in each matching specification and the complementary moderating covariates. There are some notable exceptions, however, in which we exclude or add covariates as controls. For each of the analyses in which distance to major city is the moderating covariate of interest, we exclude distance to road from the controls due to high correlation (multicollinearity). For each of the Thailand socioeconomic analyses, we add province level fixed effects to the vector of covariates. For a detailed account of the controls used in each specification see Table S6. Complete first stage results are available from the authors upon request.

Use of PLM and LOESS. We use LOESS to estimate the relationship between baseline poverty and the outcomes of interest in Costa Rica (Figure 1(a)) because we are interested in what actually happened to the poor over time rather than simply the effect of being poor. To identify the potential for protected areas to act as a mechanism for poverty traps, we do not want to partial out any of the variables that are correlated with being poor. We simply want to observe how areas with differing levels of baseline poverty fared over time.

We view the other covariates (slope, distance to city and percent agricultural workers) as moderating variables through which protection affects outcomes. For this reason we are interested in identifying the specific effect of these covariates, net of other influences, on our outcomes. Thus we use PLM. In addition, the use of PLM to isolate the specific effects of variables allows us to overlay these effects on the suitability maps with fewer concerns of confounding effects

Suitability Mapping

Motivation. The illustrative suitability maps presented in the main text characterize the suitability of end-period forested land for protection, based on past observed relationships between covariates and the environmental and socioeconomic outcomes. We characterize suitability along these two outcome dimensions because, while the targeting of protected areas is likely to be based on expected environmental outcomes, the opportunity costs of protection are socioeconomic in nature. Therefore, it would be beneficial to a planner to understand the expected joint outcomes of the establishment of protected areas.

We choose to formulate our suitability maps based on slope and distance to major city for two reasons. First, these are measurements that are globally available and have been used in past studies of protected areas. Second, these covariates capture the notion of deforestation pressure (see main text) and are therefore likely to be considered in the establishment of protected areas.

Formulation. To map expected suitability for protection in Costa Rica and Thailand we begin by rasterizing the end-period forest cover shapefiles so that each raster cell is 3 ha in size. We then create a distance to city and slope raster based on these end-period forest cover rasters for each analysis.[8]

[7] The PLM estimates were programmed in R v. 2.11.1. The code is available from the authors upon request.

[8] This leaves us with four initial rasters for each country: a distance to major city and slope raster for the deforestation analysis, and a distance to major city and slope raster for the poverty analysis.

The values of the rasters' cells are populated with measurements of distance to major city and slope, respectively.

The results from the PLM heterogeneity analyses act as the basis for our designation of expected suitability. The PLM results are appropriate for the creation of these maps because they map the continuous nonparametric *effect* of the covariates on the outcome of interest, net the effect of other influencing covariates. To allow for aggregation of suitability across covariates, we rescale the estimated covariate effects on avoided deforestation and poverty to fall within a range of 1 to 10.[9] For example, the maximum estimated effect of slope on avoided deforestation in Costa Rica is 0.139 at a slope of 14%, so it is rescaled to 10. Conversely, the minimum estimated effect is 0.00087 at a slope of 50%, so it is rescaled to 0. Similarly, all estimated effect between the `min` and `max` are rescaled and rounded. The rescaled values are then assigned to the distance to city and slope rasters for each analysis.[10] For example, all of the cells (of the slope raster) with slope values of 14 in the Costa Rica deforestation analysis are assigned a suitability score of 10. Comparable value assignments are made for each covariate in each analysis.

As a result of these assignments each parcel (in each country) has two rescaled environmental suitability scores and two rescaled socioeconomic suitability scores (one based on distance to city and the other based on slope). We use these values to calculate the average suitability (separately for environmental and socioeconomic outcomes) scores for each land parcel. Figure S1 and S2 show the aggregated environmental and socioeconomic suitability on separate maps for Costa Rica and Thailand, respectively. The final compound suitability maps (Figures 2 and 3 in the main text) are created by overlaying the aggregate environmental and socioeconomic suitability maps.

On the final suitability maps, we highlight two types of land parcels: those with expected 'win-win' outcomes (yellow), and those with expected poverty exacerbation (black). A parcel is designated as 'win-win' if its average environmental and socioeconomic suitability scores are jointly greater than or equal to 6 (this corresponds to the top five deciles). Conversely, if the underlying covariate value of a parcel is associated with negative socioeconomic impacts then the parcels is designated as unsuitable for protection due to potential poverty exacerbation from protection. For instance, due to the relationship between agricultural suitability and slope, flat parcels in Costa Rica and Thailand are designated as unsuitable for protection.

Note on Thailand Results. In the final Thailand suitability map there are distinct concentric circles of predicted 'win-win' outcomes. It can be seen from the underlying suitability maps (Figure S2) and PLM results (Figure 1(d&e)) that these expected outcomes are driven by the nonparametric relationship between the outcomes of interest and distance to a major city. Figure 1(e) indicates that the greatest poverty reduction is expected between approximately 50km and 90km. Expected avoided deforestation is also positive along this range. The range 55-75km, where both expected outcomes are relatively high, is where a majority of the 'win-win' areas lie.

While distance to major city drives the concentric circles observed in Figures 3 and S2, it is but a one facet in the determination of the joint suitability. In order for a parcel to be designated as 'win-win' there must be congruence in expected outcomes across distance to city and slope. Much of the land that lies within the 50-75km range is also relatively steeply sloped. Close examination of Figures 3 and S2 show that this is not the case throughout. In fact, there are many parcels within this range that are not designated as expected 'win-win' due to the underlying low slope.

Ancillary Analyses

Quantile Regression. In the results section of the main text, we use the LOESS estimates to assert that the establishment of protected areas has not acted as a mechanism for poverty traps in Costa Rica. Our assertion stems from the fact that, in the mapping of the LOESS, there is a general trend of greater poverty alleviation in areas with higher baseline poverty. To corroborate these results from the nonparametric LOESS estimator, we use a parametric quantile regression (see (18) for a nice overview). Quantile regressions estimate covariate effects at defined quantiles of the outcome. In our case, we use deciles of the poverty index in 2000. We are interested in the response to protection according to baseline poverty. To interpret the results of a quantile regression as a treatment effect on the distribution of outcomes, we must invoke a rank preservation assumption. This assumption implies that the poverty rank among census tracts remains stable over time. Given that the correlation coefficient between baseline and outcome poverty index is nearly 0.7, this assumption seems plausible.

We run a quantile regression (using deciles) of 2000 poverty index on an intercept and indicator of protection using the same matched set as is used in the LOESS analysis (described above). We do not include any additional controls in the regression because: (1) the LOESS estimator is (essentially) a univariate regression method, and thus our intention is to use similar specifications to that analysis; and (2) the quantile regression is run using the preprocessed matched set which is designed to be balanced across key confounding covariates. Figure S9 presents the results of the quantile regression in which the solid line represents the point estimates at each decile with the corresponding pointwise 95% confidence band in green. The point estimates can be interpreted as the effect on poverty of "moving" from unprotected to protected at each level of poverty. The results display a similar trend to that seen in the LOESS results (Figure 1(a) of the main text): namely that protection has had greater poverty alleviating effects on the poorer census tracts.

Agricultural Workers. In the main text, we use slope as a proxy for agricultural suitability. Slope has been used in a similar manner in previous studies (15) as well as a proxy for other deforestation pressures (e.g., logging access; (19)). To support the conjecture that the slope analysis is indeed highlighting the impact of opportunity costs from agriculture, we run a PLM analysis to study the heterogeneity of protection's impact conditional on baseline percentage of the workforce employed in agriculture in Costa Rica (where we have data on this measure). An opportunity cost argument would predict that avoided deforestation would be higher in areas with a high percentage of the workforce in agriculture and poverty impacts would be lower in these same areas. We observe this relationship in Figure S4 (bottom panel).

Standard Errors. All of our analyses are preceded by matching to improve balance across protected and unprotected units. Because the matching is performed with replacement there are repeated control observations in the final matched samples. The concern with repeat control observations is that

[9] Mathematica has a Rescale command which we rewrote for R.

[10] In the rescaling of the socioeconomic effects of the covariates, only positive expected outcomes are rescaled between 0 and 10. Any covariate value that is associated with socioeconomic effects deemed unsuitable for protection (see below).

precision of the standard error estimates in post-match analyses (e.g., regression) may be overstated. In response to this concern, we first note that our results are driven by the relationships presented in the ATT estimates, rather than the precision of these estimates. For example, we are more interested in the overall relationship between avoided deforestation and slope than knowing whether or not avoided deforestation was significantly different from zero at 45 percent slope.

Second, we note that the standard errors of the fit presented in the main text are not likely to be understated. The final estimate in each of our analyses (both LOESS and PLM) is designed to be interpreted in a manner similar to a post-matching, bias-adjusted difference in means. This design allows us to compare our results to the studies from which we draw. Thus we are performing the final stage LOESS using the independent variable of interest and the individual ATT, which is simply the difference between actual outcome and imputed counterfactual outcome for each protected unit (see LOESS section above). Therefore the degrees of freedom in the estimation of the standard error of the fit is based only on the number of observations in the protected sample, rather the entire sample of protected and unprotected units (as would be the case in a typical regression context). The fact that unprotected units do not add to the degrees of freedom serves to mitigate the effect of repeated observations, which lie only in the unprotected units.

Third, to offer the reader more confidence that the standard errors used in Figure 1 are not substantially understated, we calculate standard errors via bootstrapping. The 95% pointwise confidence band is determined by the 2.5 and 97.5 percentile bootstrapped outcome at each point of interest along the range of the independent variable. In each analysis, the final stage LOESS estimate is bootstrapped 1000 times.[11] The bootstrapped standard errors are overlaid on the standard errors of the fit in Figures S6-S8 in which it can be seen that the two standard error estimates coincide closely. One of the key insights that can be taken from Figures S6-S8 is that our main results are robust to alternative methods of estimating the standard errors.

Areal Interpolation

Costa Rica's census tract boundaries are not spatially consistent across time. The number of census tracts increased from 4,694 in 1973 to 17,625 in 2000. Furthermore, the addition of census tracts over time did not follow any discernible pattern, the newer subdivided census tracts do not necessarily fall within the boundaries of the old census tracts. This poses a problem for the comparability of the demographic data over time. In order to make the 2000 data comparable to the 1973 data, the geographic method of Areal Interpolation (3) is implemented, as was done in Andam et al. (2). Areal interpolation is a GIS method by which demographic variables are made comparable across time given changes in political boundaries. For our analyses, the 1973 census tracts are used as baselines. Therefore, areal interpolation assigns weights (assuming a uniform population distribution) based upon the amount that the 2000 census tracts overlap with the 1973 census tracts. These weights are used to interpolate the 2000 populations that reside within the 1973 census tract boundaries. The resulting data set contains the original 1973 demographic data according to its native boundaries and the 2000 demographic data distributed as if the census tract boundaries had not changed since 1973.

Poverty Index for Costa Rica

Ideally, a poverty measure utilized in a quasi-experimental study should be comparable across time. Costa Rica does not have properly disaggregated income data that date back to 1973 (20). Therefore, to measure the socioeconomic impacts of protected areas, an alternative metric is necessary. Principal component analysis is one method by which variables that are known to be associated with poverty can be used to form a poverty index. Other authors (4) suggest a poverty index for Costa Rica that uses indicators from the respective census to create a socioeconomic measure that is both spatially and temporally comparable. The variables included in the poverty index are (* indicates a percentage): *men in total population*, families who cook with coal or wood*, families without washing machine*, families without refrigerator*, people who are employed and get a salary as job remuneration*, illiterate population aged 12 or more*, household dwellings without connection to private or public water system*, household dwellings without sewers*, household dwellings without electricity*, household dwellings without telephone*, dwellings with earth floor*, dwellings in bad condition*, dwellings without bathroom*, dwellings without access to hot water*, dependency ratio, average number of occupants per bedroom, average years of education per adult.* A similar measure was employed by the Mexican government in the analysis of the PROGRESA program (4).

The poverty index for 1973 is used as a covariate for preprocessing prior to all socioeconomic and deforestation analyses. The 2000 poverty index is used as the outcome for all of the socioeconomic analyses. See Andam et al. (2) for a more detailed explanation of the poverty index estimation.

1. Andam KS, Ferraro PJ, Pfaff A, Sanchez-Azofeifa G, & Robalino JA (2008) Measuring the effectiveness of protected area networks in reducing deforestation. Proc. Natl. Acad. Sci. U.S.A. 105(42):16089-16094.
2. Andam K, Ferraro P, Sims K, Healy A, & Holland M (2010) Protected areas reduced poverty in Costa Rica and Thailand. Proc. Natl. Acad. Sci. U.S.A. 107(22):9996.
3. Reibel M (2007) Geographic information systems and spatial data processing in demography: A review. Population Res. Pol. Rev. 26(5):601-618.
4. Cavatassi R, Davis B, & Lipper L (2004) Estimating poverty over time and space: construction of a time-variant poverty index for Costa Rica. (ESA Working Paper).
5. Sims KRE (2010) Conservation and Development: Evidence from Thai Protected Areas. J. Environ. Econ. Manage. 60: 94-114.
6. Healy A & Jitsuchon S (2007) Finding the poor in Thailand. J. Asian Econ. 18(5):739-759.
7. Elbers C, Lanjouw J, & Lanjouw P (2003) Micro-level estimation of poverty and inequality. Econometrica:355-364.
8. Ho D, Imai K, King G, & Stuart E (2007) Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. Polit. Anal. 15:199-236.
9. Imbens G & Wooldridge J (2009) Recent developments in the econometrics of program evaluation. J. Econ. Lit. 47(1):5-86.
10. Holland P (1986) Statistics and causal inference. J. Amer. Statistical Assoc. 81(396):945-960.
11. Angrist J & Pischke J (2009) Mostly harmless econometrics: an empiricist's companion (Princeton Univ Pr).
12. Abadie A, Drukker D, Herr J, & Imbens G (2004) Implementing matching estimators for average treatment effects in Stata. SJ 4:290-311.
13. Cleveland W (1979) Robust locally weighted regression and smoothing scatterplots. J. Amer. Statistical Assoc.:829-836.
14. Cleveland W & Devlin S (1988) Locally weighted regression: an approach to regression analysis by local fitting. J. Amer. Statistical Assoc. 83(403):596-610.
15. Ferraro PJ & Hanauer MM (2010) Protecting Ecosystems and Alleviating Poverty with Parks and Reserves: 'Win-Win' or Tradeoffs? Environmental and Resource Economics. In Press.
16. Yatchew A (1997) An elementary estimator of the partial linear model. Econ. Letters 57(2):135-143.

[11] The bootstrapping function was written in R v. 2.11.1. Code is available from authors upon request.

17. Yatchew A (1998) Nonparametric regression techniques in economics. J. Econ. Lit. 36(2):669-721.
18. Lokshin M (2006) Difference-based semiparametric estimation of partial linear regression models. SJ 6(3):377.
19. Koenker R & Hallock K (2001) Quantile regression. Journal of Economic Perspectives 15(4):143-156.
20. Gindling T & Terrell K (2004) Legal minimum wages and the wages of formal and informal sector workers in Costa Rica. in William Davidson Institute Working Papers Series (University of Michigan).

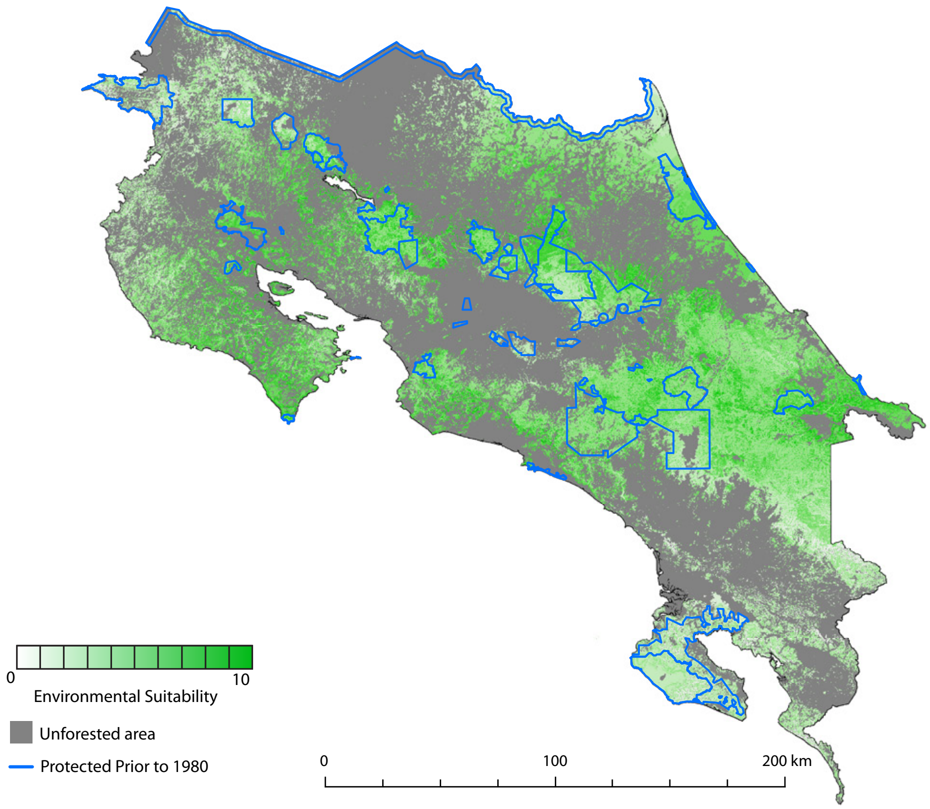

Figure S1a. Costa Rica Environmental Suitability Map

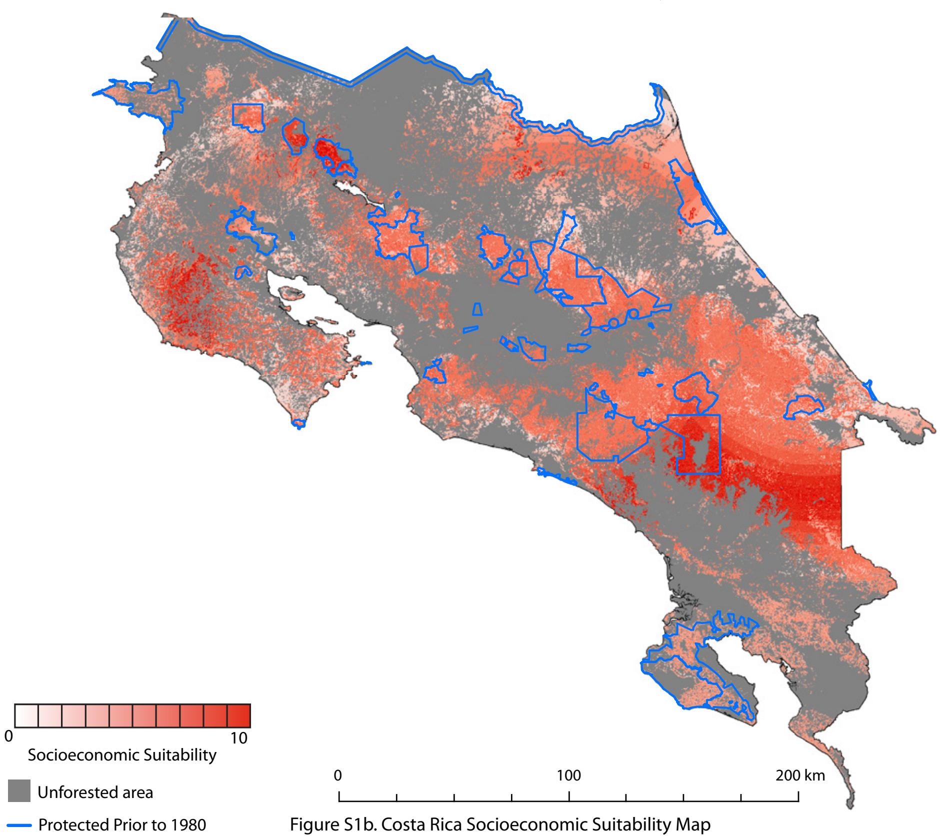

Figure S1b. Costa Rica Socioeconomic Suitability Map

Fig.S1. Costa Rica protected area suitability maps by environmental and socioeconomic suitability.

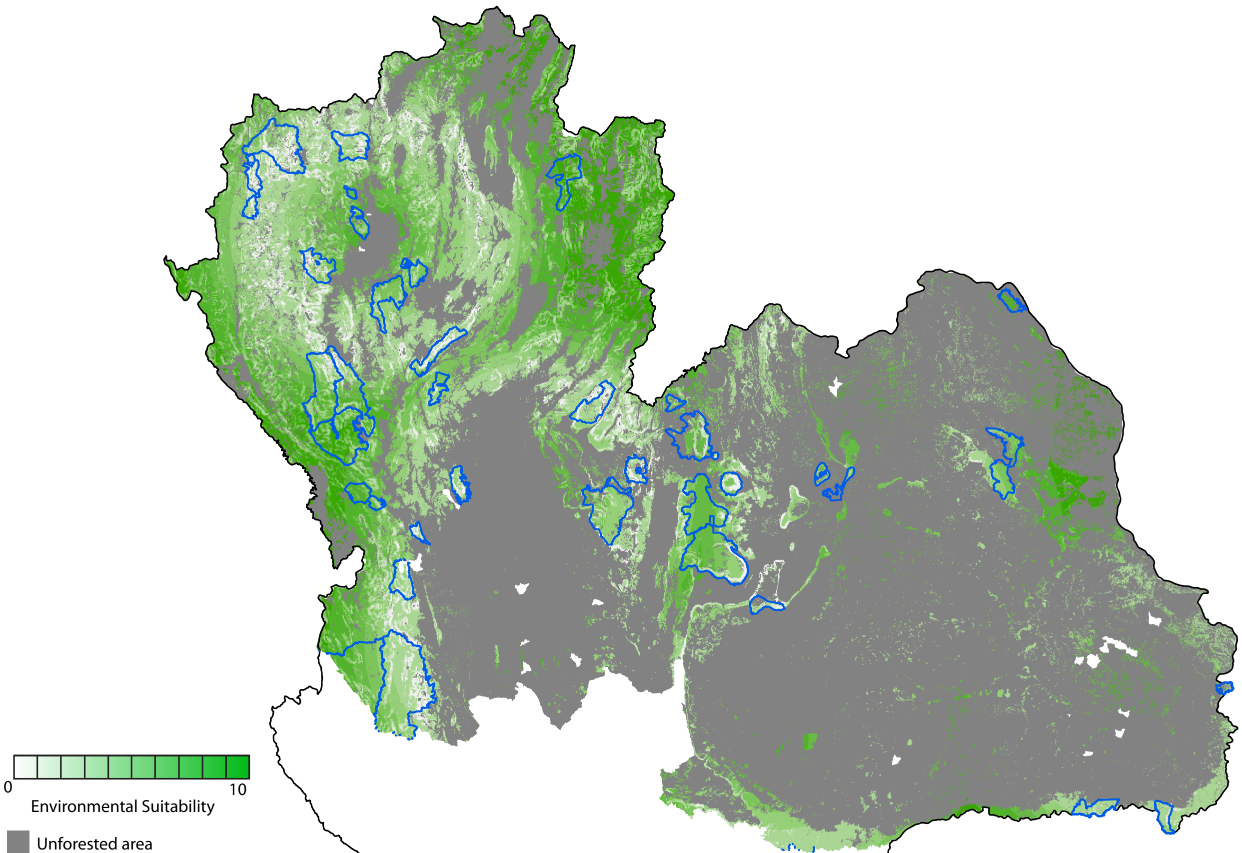

Figure S2a. Thailand Environmental Suitability Map

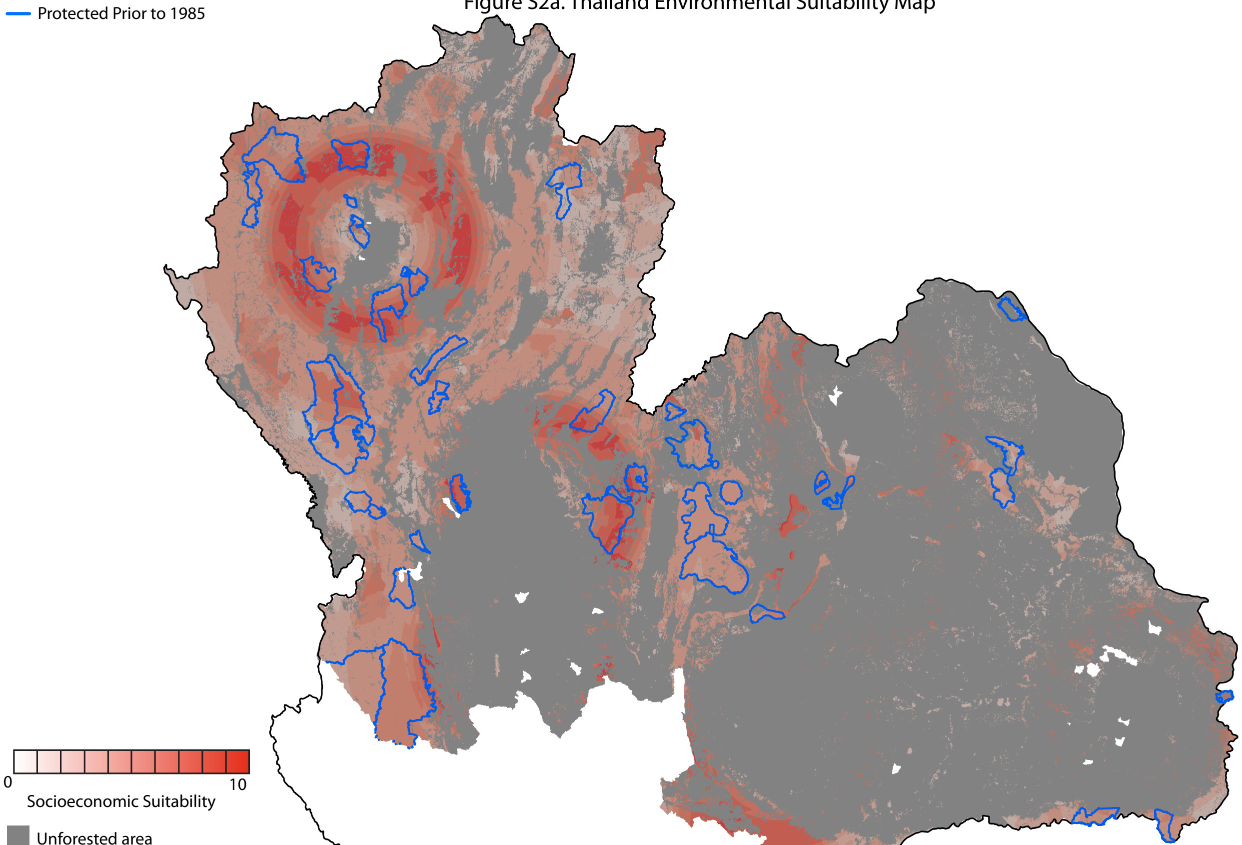

Figure S2b. Thailand Socioeconomic Suitability Map

Fig.S2. Thailand protected area suitability maps by environmental and socioeconomic suitability.

Fig.S3. Costa Rica: Full LOESS results.

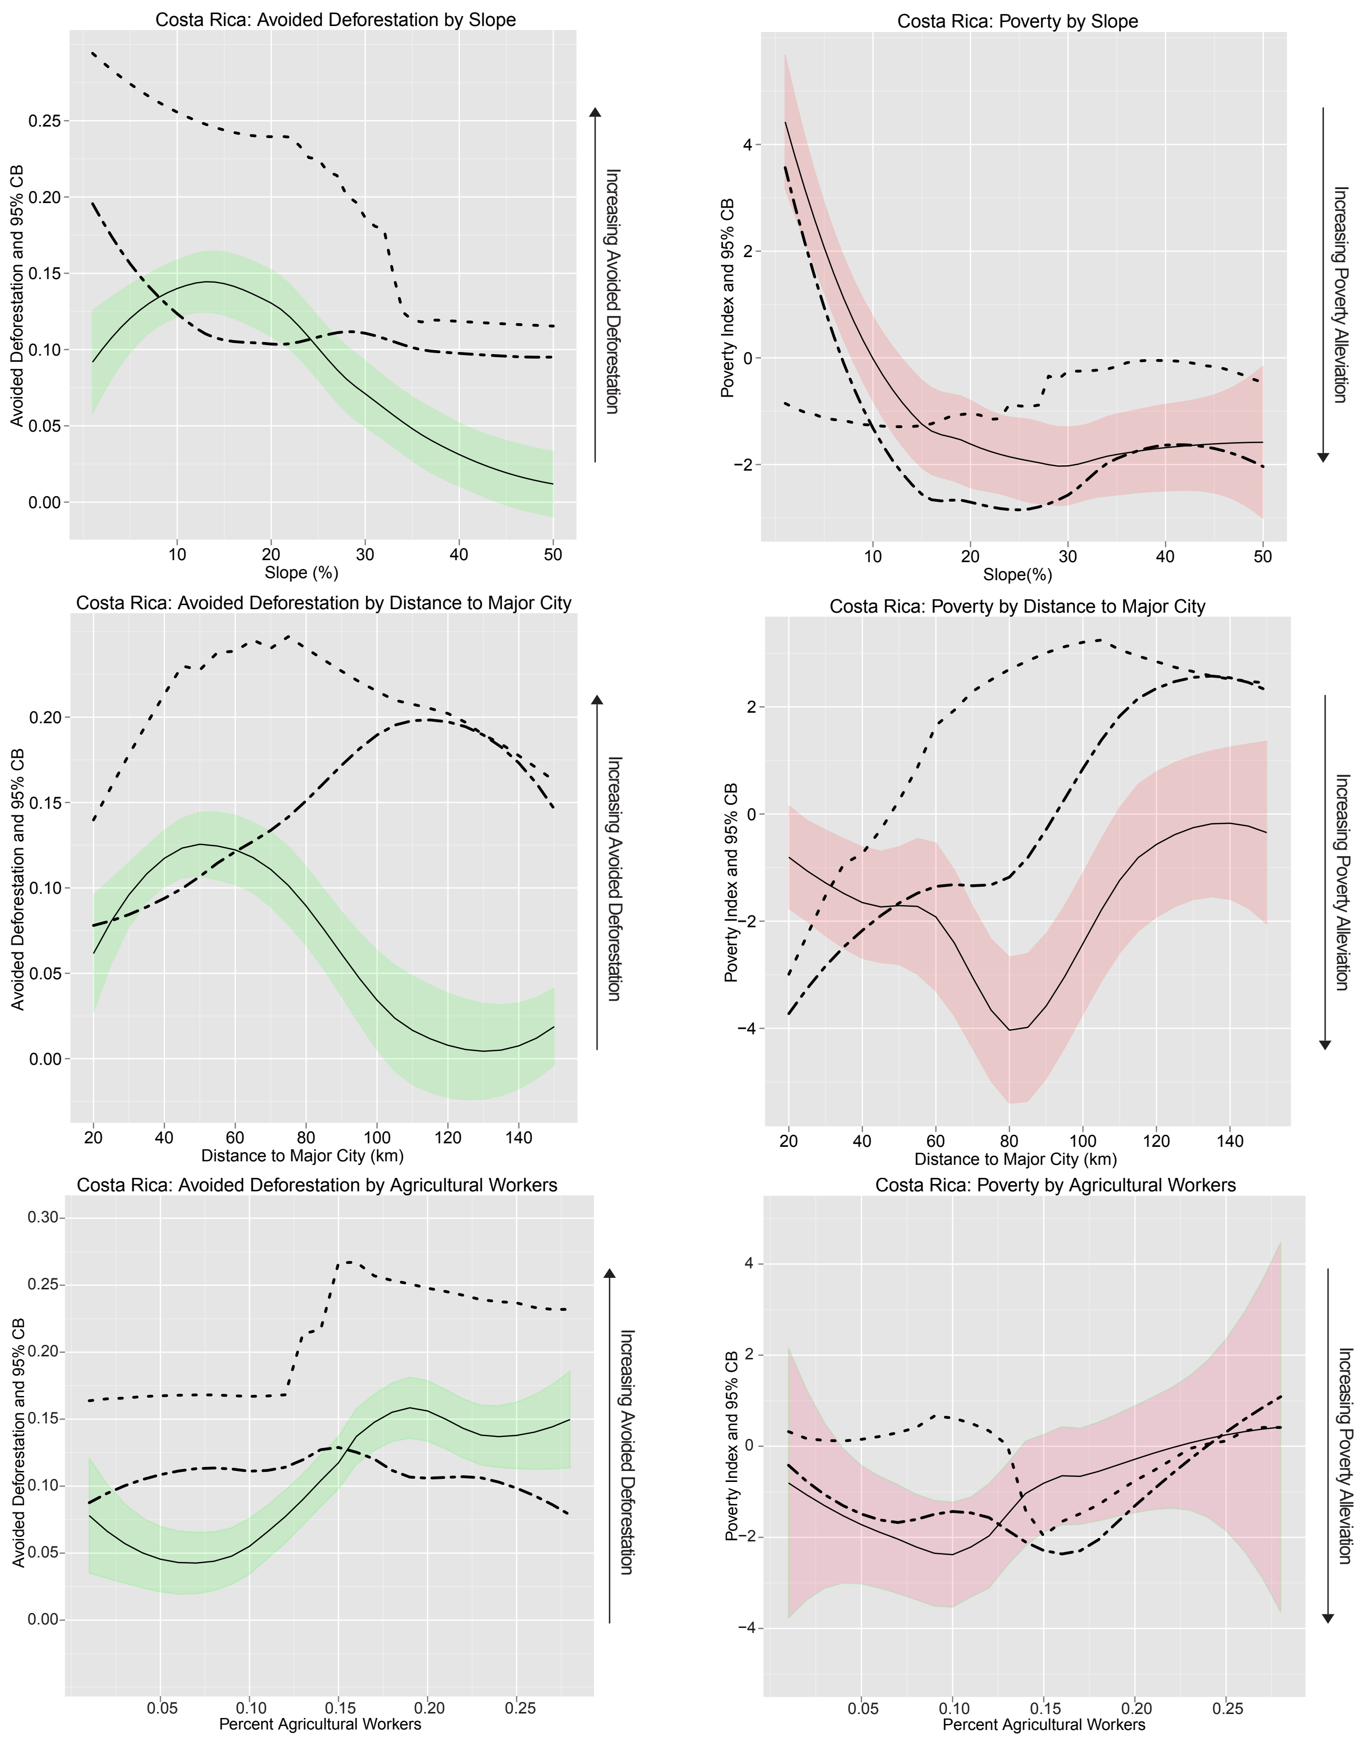

Fig.S4. Costa Rica: full heterogeneous response to protection results.

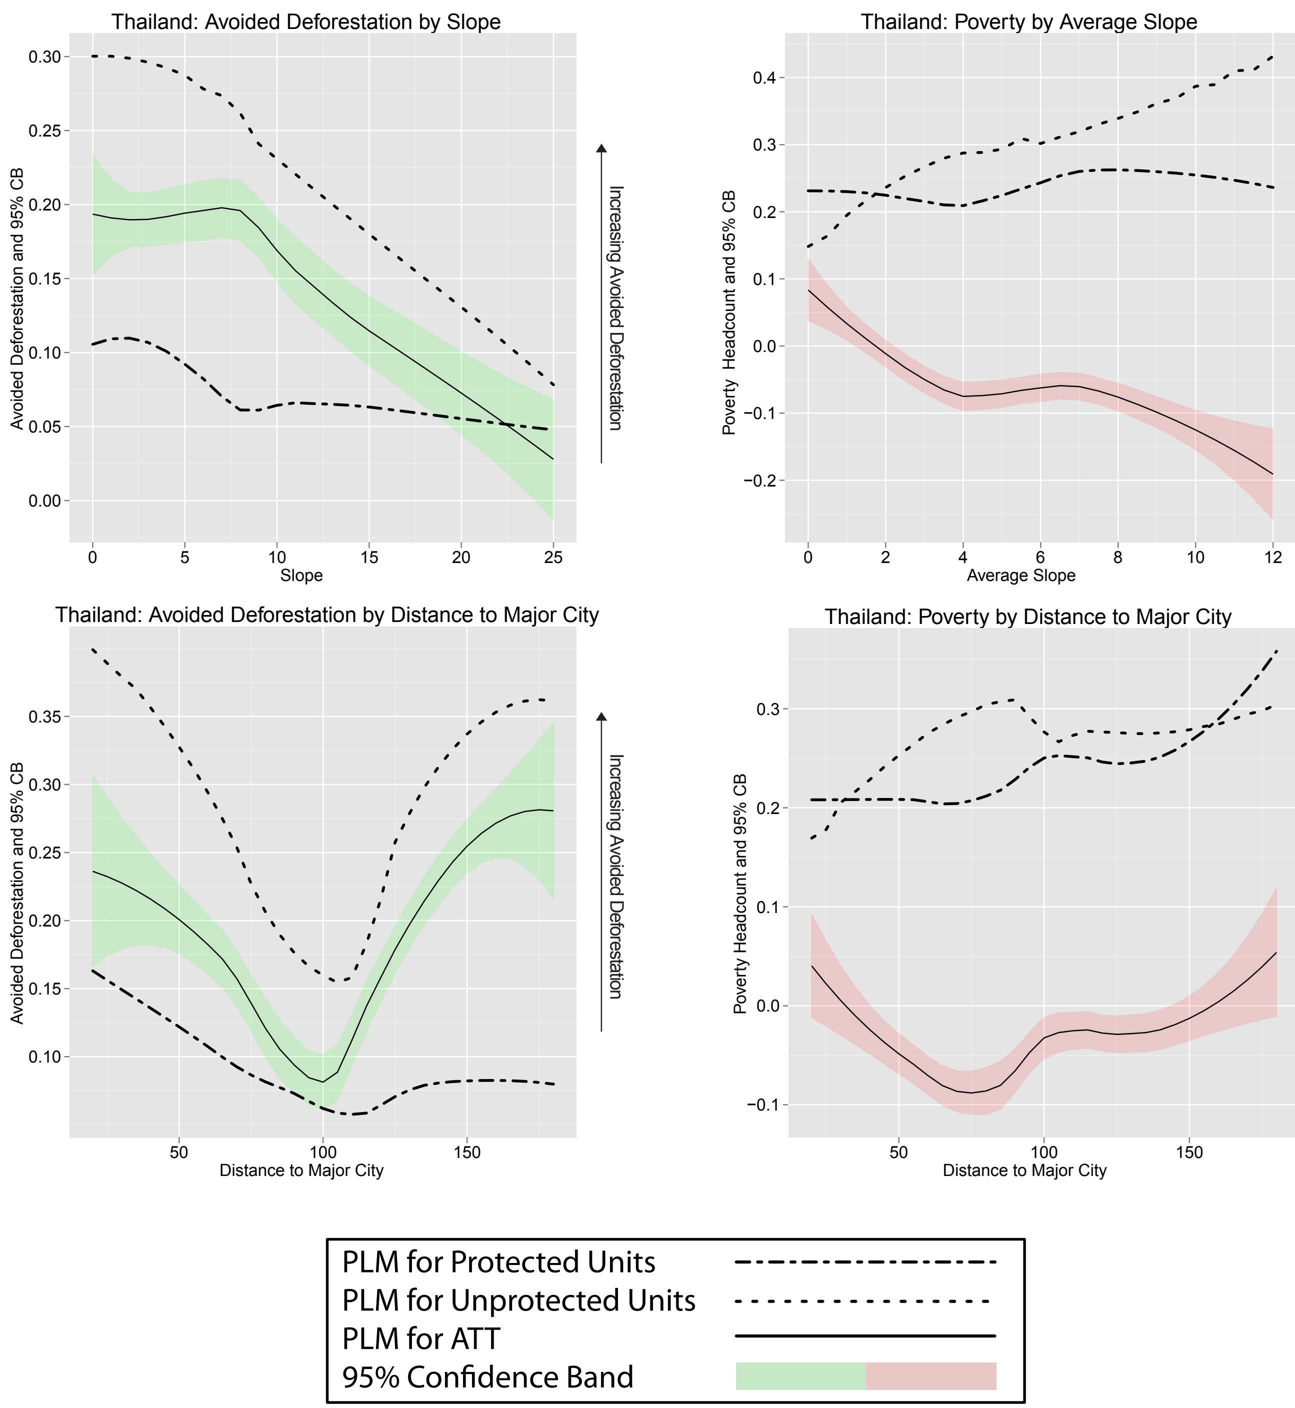

Fig.S5. Thailand: full heterogeneous response to protection results.

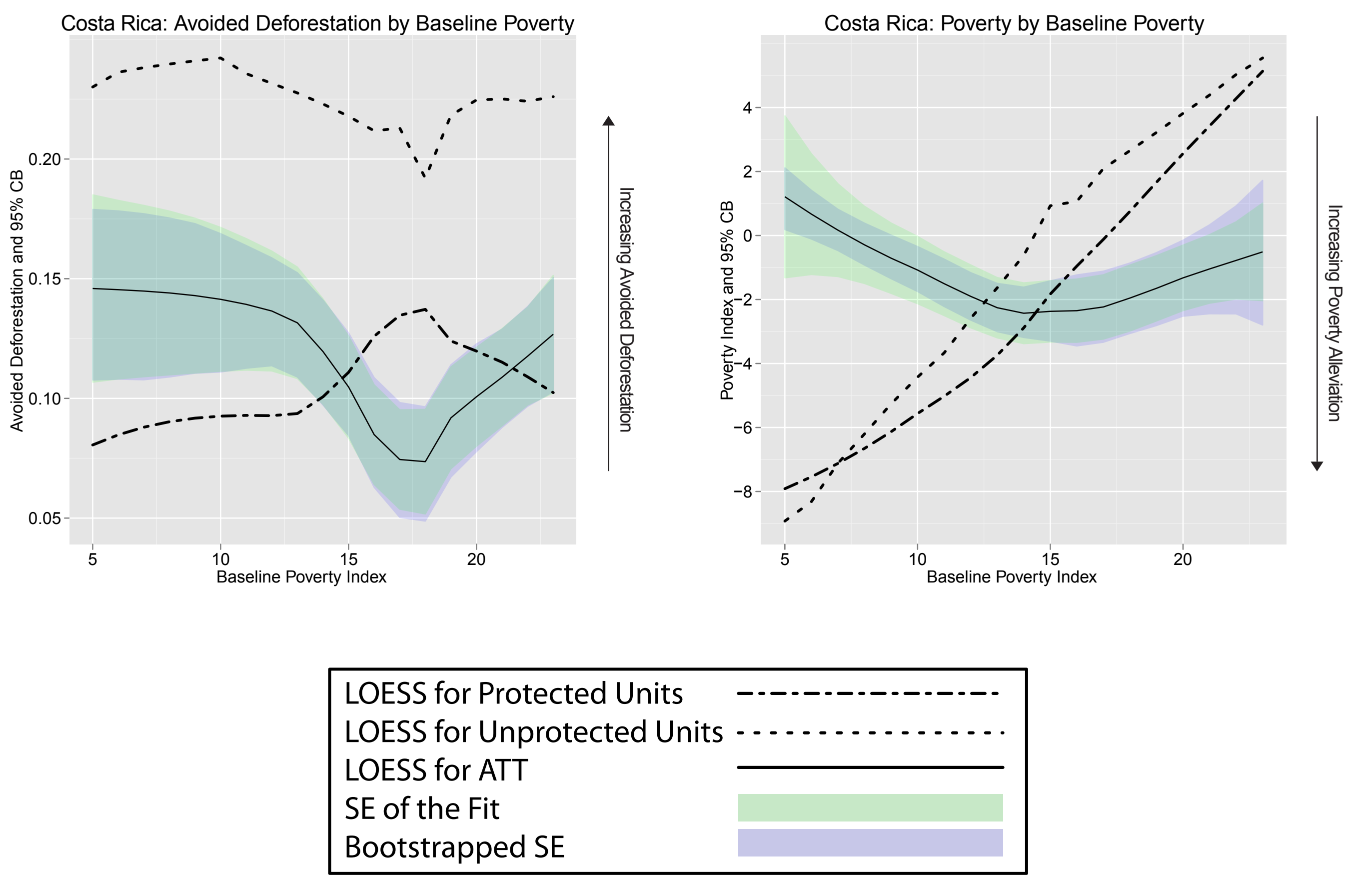

Fig.S6. Costa Rica: Comparison of bootstrapped standard errors to standard errors of the fit.

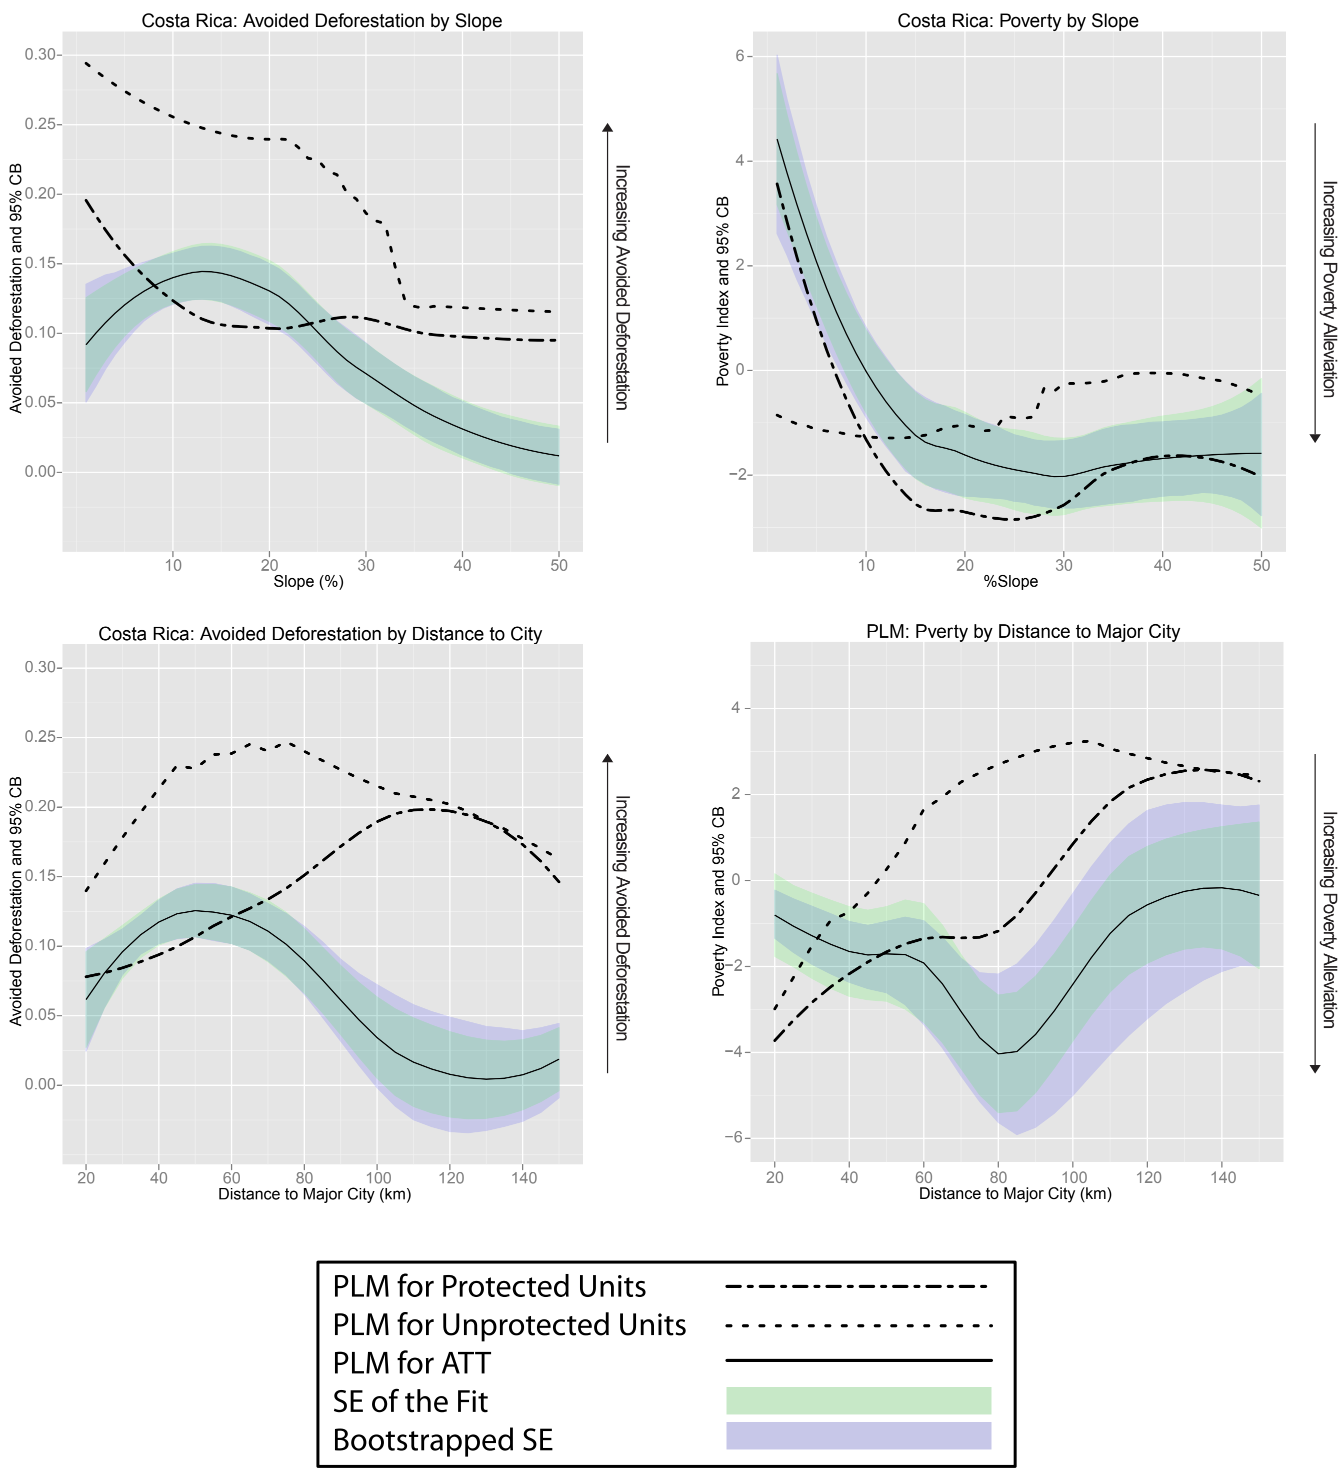

Fig.S7. Costa Rica: Comparison of bootstrapped standard errors to standard errors of the fit.

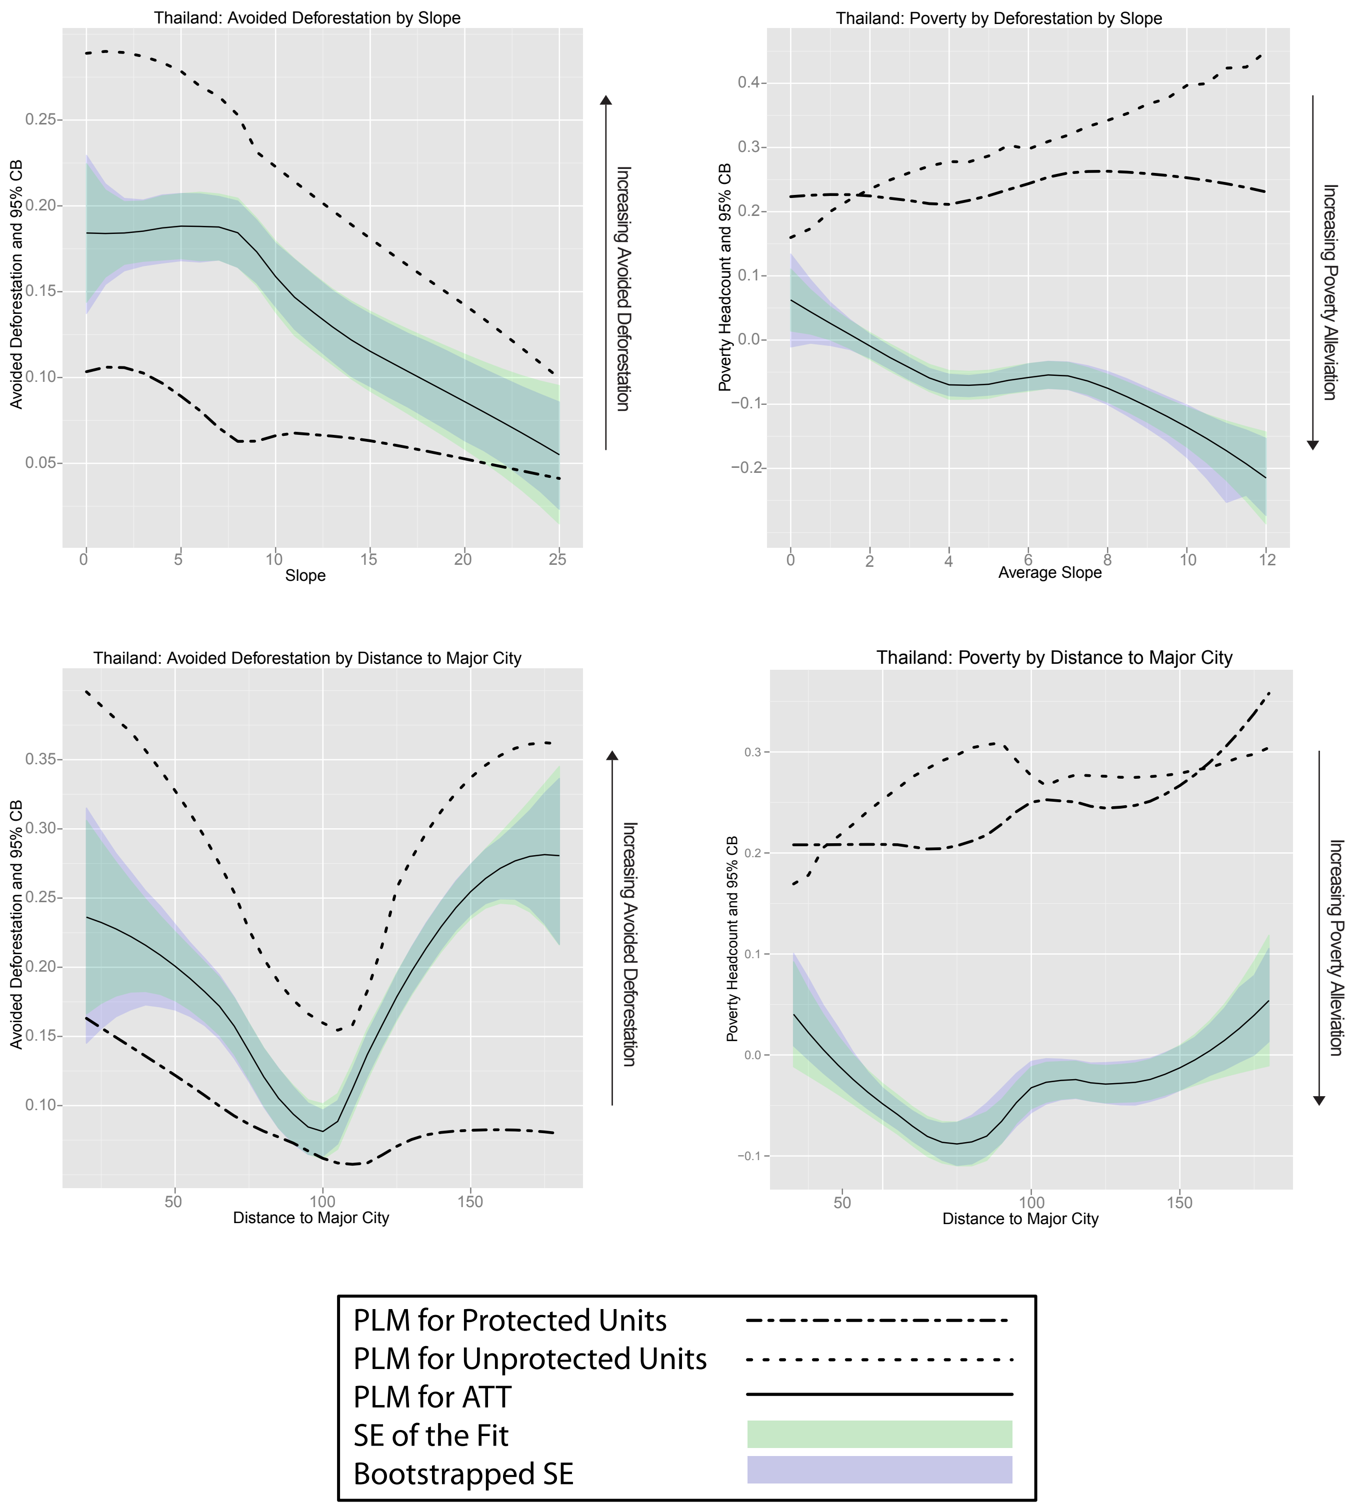

Fig.S8. Thailand: Comparison of bootstrapped standard errors to standard errors of the fit.

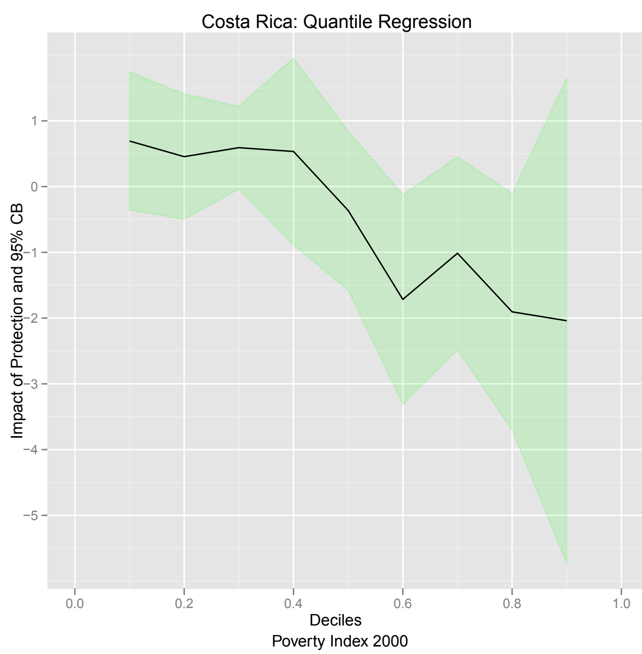

Fig.S9. Costa Rica: Quantile regression estimating impact of protection according to deciles of 2000 poverty index.

	Variable	Description	Mean	Median	Standard Deviation	Range
Deforestation Covariates	High Productivity Land	Land Use Capacity I, II & III Land suitable for agricultural production. May require special land and crop management (classes II & III).	0.008	0	0.09	0-1
	Medium-High Productivity Land	Land Use Capacity IV Moderately suitable for agricultural production; permanent of semi-permanent crops	0.0289	0	0.167	0-1
	Medium-Low Productivity Land	Land Use Capacity V, VI & VII Strong limiting factors on agricultural production.	0.0802	0	0.272	0-1
	Distance to Forest Edge	Distance (km) to the edge of the forest in 1960	2.79	2.35	2.19	0.0001-11.2
	Distance to Road	Distance (km) to nearest road in 1969.	16.99	14.28	11.62	0.04-53.31
	Distance to Major City	Distance (km) to nearest major city: Limon, Puntarenas or San Jose.	77.4	56.9	49.53	9-180.5
Socioeconomic Covariates	Baseline Poverty	Poverty index measured in 1973.	14.9	15.8	6.43	-6.4-28.9
	Forest Cover	Percentage of census tract occupied by forest in 1960.	0.412	0.383	0.342	0-1
	% High Productivity Land	Percent of census tract occupied by Land Use Capacity I, II & III land.	0.118	0	0.22	0-1
	%Medium-High Productivity Land	Percent of census tract occupied by Land Use Capacity IV land.	0.295	0.04	0.377	0-1
	%Medium-Low Productivity Land	Percent of census tract occupied by Land Use Capacity VI, VII or VIII land.	0.347	0.156	0.387	0-1
	Distance to Major City	Average distance (km) from each $300m^2$ land plot within a census tract to nearest major city: Limon, Puntarenas or San Jose.	57.3	49.7	41.28	0.0037-208
	Roadless Volume	The sum of the product of area and distance to nearest road (1969) for every square with side length 100m within the census tract.	308,000	66,400	699,100	0.28-7,590,000

Table S1. Costa Rica - Summary statistics and description of covariates used as controls to form counterfactual samples.

	Variable	Description	Mean	Median	Standard Deviation	Range
Deforestation Covariates	Slope	Slope of parcel (degrees)	5.905	5	5.48	0-43
	Distance to Major River	Distance (km) to major river (flow accumulation > 5000)	30.598	27.62	19.552	0.004-109.3
	Elevation	Elevation (m) of parcel	555.535	497	316.942	0-2183
	Distance to Forest Edge	Distance (km) to the edge of the forest in 1973	2.747	1.884	2.775	0.0001-19.58
	Distance to Road	Distance (km) to nearest road in 1962	21.08	16.4	17.682	0.00076-93.8
	Distance to Major City	Distance (km) to nearest major city (pop > 100,000)	113.573	113.5	42.621	7.26-254.3
Socioeconomic Covariates	Average Slope	Average slope of subdistrict (degrees)	1.018	0.0504	2.042	0-14.33
	Maximum Slope	Maximum slope of subdistrict (degrees)	4.05	0.9882	6.99	0-46.99
	Distance to Major River	Distance (km) to major river (flow accumulation > 5000)	21.61	0	16.61	0.01-97.82
	Forest Cover 1973	Percent of subdistrict covered by forest, 1973	0.194	0.00423	0.315	0-1
	Distance to Major City	Distance (km) to nearest major city (pop > 100,000)	85.59	81.03	44.51	10.05-222.6
	Distance to Major Road	Distance (km) to major road in 1962	5.26	7.615	6.22	0.002-76.16
	Distance to Any Road	Distance (km) to minor road in 1962	10.42	3.448	0.002	88.08
	Distance to Thai Border	Distance (km) to Thailand border	91.62	91.33	52.36	0.062-218.9
	Near Watershed	Within 1 km of major watershed boundary	0.461	0	0.499	0-1
	Distance to Rail Line	Distance (km) to rail line	55.05	42.95	45.76	0.015-222.1
	Dist. to Mineral Deposit	Distance (km) to nearest mineral deposit	119.46	102.7	84.73	1.371-376.4
	Temperature	Average temperature (°C) for subdistrict	25.37	25.89	1.448	18.07-27.85
	Rainfall	Average monthly rainfall (mm)	1064	1021	225.3	375.8-2308

Table S2. Thailand - Summary statistics and description of covariates used as controls to form counterfactual samples.

Covariate	Status	Mean Prote-cted Plots	Mean Cont-rol Plots	Difference in Mean	Normalized Difference	Mean eQQ Difference	% Improve Mean Diff.
High Land Use Capacity	Unmatched	0.008	0.205	-0.197	0.307	0.197	
	Matched	0.008	0.008	0.000	0.000	0.000	100.0%
Medium-High Land Use Capacity	Unmatched	0.029	0.198	-0.170	0.259	0.170	
	Matched	0.029	0.029	0.000	0.000	0.000	100.0%
Medium-Low Land Use Capacity	Unmatched	0.080	0.507	-0.427	0.563	0.427	
	Matched	0.080	0.080	0.000	0.000	0.000	100.0%
Distance to Forest Edge	Unmatched	2.857	2.045	0.812	0.162	0.886	
	Matched	2.857	2.713	0.143	0.031	0.148	82.3%
Distance to Road	Unmatched	17.354	15.336	2.017	0.078	2.099	
	Matched	17.354	16.709	0.645	0.026	0.975	68.0%
Distance to Major City	Unmatched	76.980	80.515	-3.535	0.037	15.894	
	Matched	76.980	77.912	-0.933	0.008	2.295	73.6%

Table S3. Costa Rica - Covariate balance for baseline avoided deforestation analysis.

	Difference in Means	Mahalanobis Matching†
Avoided Deforestation ($Y_{protected}$ - $Y_{unprotected}$)	-0.2595*** {0.0062}	-0.14738*** (0.0175)
N Treated	NA	2,808
(N Available Controls)	NA	(13,609)

*** Indicates significance at the 1% level

† ATT is post-match difference in means using regression bias adjustment to control for bias in finite samples

(Abadie-Imbens heteroskedasticity robust standard errors)

{Standard errors}

Table S4. Thailand - Baseline avoided deforestation analysis.

Covariate	Status	Mean Protected Plots	Mean Control Plots	Difference in Mean	Normalized Difference	Mean eQQ Difference	% Improve Mean Diff.
Distance to Major City	Unmatched	109.670	114.061	-4.395	0.040	8.498	
	Matched	109.670	110.422	-0.756	0.008	4.958	82.8%
Distance to Road	Unmatched	31.160	17.161	13.999	0.401	14.001	
	Matched	31.160	29.419	1.741	0.039	2.054	87.6%
Distance to Forest Edge	Unmatched	3.620	2.261	1.359	0.242	1.358	
	Matched	3.620	3.304	0.316	0.051	0.317	76.7%
Slope	Unmatched	7.960	4.989	2.970	0.254	2.972	
	Matched	7.960	7.807	0.152	0.012	0.437	94.9%
Distance to Major River	Unmatched	35.930	27.237	8.691	0.217	8.689	
	Matched	35.930	34.885	1.043	0.024	2.464	88.0%
Elevation	Unmatched	697.130	486.612	210.519	0.307	210.448	
	Matched	697.130	635.127	62.004	0.093	62.061	70.6%

Table S5. Thailand - Covariate balance for baseline avoided deforestation analysis.

Covariate	Exclusions	Inclusions	Justification
Slope	Land Use Capacity[†‡]	NA	LUC is a function of slope
Distance to Major City	Distance to Road[†‡]	NA	Colinearity with distance to city
% Agricultural Workers	NA	NA	NA
Slope	NA	Province Fixed Effects[‡]	Control for baseline poverty
Distance to Major City	Distance to Road[†‡]	--	Colinearity with distance to city
	Distance to Railroad[‡]	--	Colinearity with distance to city
	--	Province Fixed Effects[‡]	Control for baseline poverty

Costa Rica Thailand

Baseline set of controls for each analysis include all matching covariates (Table S1) and other mediating covariates

[†] Indicates exclusion/inclusion from the deforestation analysis

[‡] Indicates exclusion/inclusion from the socioeconomic analysis

Table S6. Exclusions and Inclusions, with respect to the baseline set of controls, in PLM analyses.